

Fabrice Marchal, Kai Nagel

Modeling Location Choice of Secondary Activities with a Social Network of Cooperative Agents

Journal article | Accepted manuscript (Postprint)

This version is available at <https://doi.org/10.14279/depositonnce-7801>



Fabrice Marchal, Kai Nagel, Modeling Location Choice of Secondary Activities with a Social Network of Cooperative Agents, Transportation Research Record: Journal of the Transportation Research Board (Vol 1935) pp. 141-146. Copyright © [2005] (Sage). DOI: [10.3141/1935-16].

Terms of Use

Copyright applies. A non-exclusive, non-transferable and limited right to use is granted. This document is intended solely for personal, non-commercial use.

WISSEN IM ZENTRUM
UNIVERSITÄTSBIBLIOTHEK

Technische
Universität
Berlin

Modeling Location Choice of Secondary Activities with a Social Network of Cooperative Agents

F. Marchal and K. Nagel

Activity-based models in transportation science focus on the description of human trips and activities. Modeling the spatial decision for so-called secondary activities is addressed in this paper. Given both home and work locations, where do individuals perform activities such as shopping and leisure? Simulation of these decisions using random utility models requires a full enumeration of possible outcomes. For large data sets, it becomes computationally unfeasible because of the combinatorial complexity. To overcome that limitation, a model is proposed in which agents have limited, accurate information about a small subset of the overall spatial environment. Agents are interconnected by a social network through which they can exchange information. This approach has several advantages compared with the explicit simulation of a standard random utility model: (a) it computes plausible choice sets in reasonable computing times, (b) it can be extended easily to integrate further empirical evidence about travel behavior, and (c) it provides a useful framework to study the propagation of any newly available information. This paper emphasizes the computational efficiency of the approach for real-world examples.

Activity-based models in transportation science focus on the description of the organization of human activities in time and space. This organization determines the demand for travel, that is, the number of users that the various transportation systems need to accommodate. It is assumed that the demand for travel is derived from the demand for performing activities at specific locations. Obviously, individuals constantly perform some trade-off between enjoying activities that have a high reward value (for instance, working at a company in the downtown area) and the time and budget it takes to reach the specific location of these activities. Various operational models such as URBANSIM are available to describe this trade-off for the choice of home and work locations (1). It is essentially assumed that users perform a trade-off between rents, travel costs, and wages. However, empirical evidence has shown that a significant amount of traffic is generated for purposes other than commuting, often referred to as secondary activities: shopping, leisure, going to social events, and so forth (2).

F. Marchal, Swiss Federal Institute of Technology, IVT ETH Hönggerberg, CH-8093 Zurich, Switzerland. K. Nagel, Institute for Land and Sea Transport Systems, Technical University of Berlin, Sek SG 12, Salzufer 17-19, D-10587 Berlin, Germany.

PROBLEM STATEMENT

This work intends to model the specific process of the location choice of secondary activities in the case of high-resolution data sets. The methodological constraints are that the modeling should be behaviorally sound, compatible with microeconomics foundations, and computationally feasible. The temporal dimension (i.e., the scheduling of activities) is ignored for the time being. It is assumed that the order of activities, called a plan, is given (i.e., leaving home, going to work, working for 8 h, going shopping at lunch time, etc.). The physical environment is described by two large data sets that typically originate from geographical information systems (GIS): (a) land-use data and (b) transportation system data. The land use is a raster-type description that includes the information about the nature of each parcel of the studied area (e.g., housing density, number of shops, and type of area—rural, commercial, industrial). The transportation system is a vector-type description of the various transportation modes available (e.g., car, rail, bus) as a network with nodes and links. Today, these data are available at a very high resolution: typical land-use cells are 100 m², and road networks are described down to 10-m road sections. The long-term goal of this research is to model entire metropolitan areas microscopically by simulating the individual decisions of millions of citizens. Therefore, the problem at hand can be stated as follows: how to simulate the selection of the activity locations of $A = 10^6$ citizens in a grid that has $C = 10^5$ cells. Note that the travel times from cell to cell have to be given by external traffic models. Multiagent traffic assignment models such as those developed by Raney et al. and de Palma and Marchal are now able to predict travel time patterns for large-scale data sets (3, 4). Therefore, that issue will not be considered. The generation of travel demand for traffic models is addressed here.

MICROECONOMICS FOUNDATION

In transportation science the standard practice to approach such problems is to use random utility models (RUMs) borrowed from the discrete choice theory of microeconomics (5, 6). These models assume that individuals are maximizing their own utility. For instance, the utility to go shopping at a mall located in cell i for a simple plan (i.e., home–shopping–work) is given by the following:

$$U_i = R_i - C_{hi} - C_{iw} + \mu \epsilon_i = V_i + \mu \epsilon_i$$

where

R_i = reward associated with shopping at that particular facility, which depends on the availability of goods, their prices, and so on;

C_{hi} = travel cost to travel from home to cell i ;

C_{iw} = travel cost to travel from cell i to work;
 μ = scale factor; and
 ϵ_i = random variable that is specific to the individual.

The latter random utility part captures all the hidden preferences of a specific user for location i that are not accessible to the modeler. By contrast, V_i is called the deterministic part of the utility. Under the assumption that ϵ_i is i.i.d. extreme value distribution of Type I, it can be shown that the probability to choose to go shopping at cell k is given by the following:

$$P(k) = P(U_i < U_k \forall i \neq k) = \frac{\exp(-V_k/\mu)}{\sum_{i=1}^c \exp(-V_i/\mu)}$$

Because a probability greater than zero is assigned to each potential intermediary stop on a cell, this formulation requires a full enumeration of the possibilities on the spatial grid. For trips that count S intermediary stops, the complexity is $O(AC^S)$, which is not feasible in realistic cases. Initially, RUMs were intended for the description of choices between a small set of alternatives distinguishable by humans (e.g., car brands). Their application to a discretized continuum (i.e., urban space) remains behaviorally questionable. Still, it is believed that it is fundamental to keep some compatibility with RUMs because of the literature that has been devoted to developing empirical techniques (e.g., surveys) to calibrate the parameters of those models.

Another drawback of RUMs is that they provide only a static representation that does not take into account the temporal dimension of the decision process. RUMs do not model explicitly the choice process but only its outcome. The goal is to develop a model that reflects the underlying learning process and that will eventually include the dynamics of the building of the choice set. Capturing that aspect is potentially important to study, for instance, the evolution of travel demand given some modification in land-use patterns. Moreover, the environment itself should be time-dependent. Travel times and travel impedances to move in the network are subject to within-day and daily variability. The feedback of the travel conditions on travel choices and conversely is also often missing from static RUM-based analysis. To overcome these limitations, a four-phase multiagent simulation model is proposed that includes a dynamic learning process and that can, in principle, be coupled with a dynamic mobility simulator. Agents travel, explore, learn, and socialize.

MULTIAGENT-BASED APPROACH

A multiagent-based simulation is proposed in which each agent (i.e., each simulated citizen) has only limited, accurate information about N cells ($N \ll C$), called the “memory” of the agent. The intuition is that real humans have limited cognitive abilities and can consider only a small number of options at the same time. The organization of these options in the mind and the human representation of space is probably far different from “pixels.” Nevertheless, that trivial representation is kept for now; it can be replaced later by a more sophisticated one, for instance, that of Arentze and Timmermans (7).

It is assumed that agents are interconnected by a social network through which they can exchange information about their respective subsets. Each agent is socially connected to K acquaintances or “kins.” The intuitive benefit of the social network is that agents are very heterogeneous yet face similar choices. Therefore, a decision

that is optimal for an agent might be close to optimal for one of his or her acquaintances. It is shown below that the diffusion of the knowledge of optimal strategies through the social network can, indeed, exploit that hidden redundancy.

The simulation is iterative, and each round has four stages: evaluation, socialization, exchange, and exploration.

Evaluation

Each agent performs the location choice of the intermediary stops based on his or her own private information. The choice can be deterministic (the best cells are selected from the memory of the agent) or probabilistic (a RUM is applied to the small set corresponding to the memory of the agent). The computing load for building and storing the travel plans is $O(AN^S)$. At this point, the plans are fed in a dynamic traffic model [such as MATSIM (3) or METROPOLIS (4)]. The traffic simulator computes the delays incurred due to traffic congestion, which are then used in the next round of evaluation.

Socialization

Social connections are created and deleted dynamically. The deletion mechanism is a simple exponential decay. The creation mechanism is a spatial reinforcement reminiscent of pheromones in ant colonies optimization (8).

Exchange

For each social connection, an agent has the opportunity to exchange a piece of information. A cell is picked randomly from the agent’s memory, and the other agent is informed about it. The exchange is bidirectional, and the outcome of the exchange is described by the learning mechanism below. More sophisticated exchange strategies could be taken into account.

Exploration

Agents have the possibility of exploring cells in the neighborhood of those that they visit. This stage is intended mainly to recover potential information lost in the other stages, thus relieving the implementation from checking that any cell was lost from the global knowledge of all the agents. Obviously $O(A)$ operations are required.

LEARNING MECHANISM

The memory of an agent is represented in Figure 1: the first two buffers contain the information about locations that are either close to home or close to work (e.g., a small circular area). The third buffer called the “elite” buffer (of size E) corresponds to locations that have high score values, and the last buffer contains “vague” information about cells that have poor score values. The scores corresponding to the elite buffer are also kept in the agent’s memory. When an agent informs another agent during the exchange stage, a cell is picked randomly from the total memory of the informer. At that point, the informed agent evaluates how this new cell information can potentially improve his or her plan score. That is done by comparing the

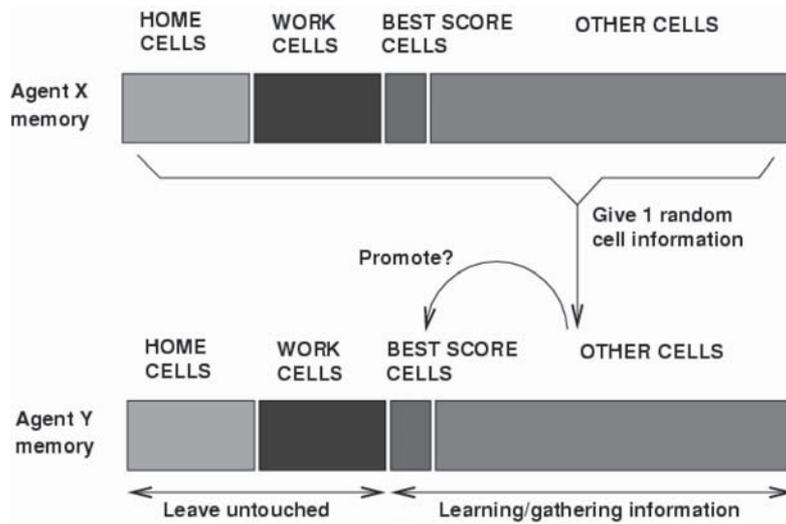


FIGURE 1 Agent memory organization and learning.

poorest score of the elite buffer to the scores of all the potential plans with at least one stop at the new cell. This revision implies replacing each intermediary stop of the potential plans by the new cell: $O(SN^{(S-1)})$ operations are needed. If the score is better than the worst solution of the elite buffer, the new cell is promoted to that buffer and it is sorted: $O(\ln E)$ operations are needed. If the cell does not improve one of the elite plans, the cell information replaces a cell randomly selected from the vague buffer. That has two consequences. First, agents keep information that is not relevant to themselves but that might be to others in the future, hence they adopt a cooperative behavior that is not supervised. Second, the information in the vague buffer can be erased and lost forever. The information about cells with low utility is more likely to be lost from the collective memory. The exploration phase allows them to still be recovered. So far, the computation load of a single round is $O(A(\ln E + SN^{(S-1)}))$, which is feasible for reasonable assumptions ($S \leq 3$, $N < 50$, $E < 50$). However, the number of iterations is still to be determined. Note that the learning speed of the overall process depends on the greediness of the exchange, which is a function of the ratio of the size of the elite buffer and the size of the vague buffer.

SOCIAL NETWORK DYNAMICS

Social links disappear following an exponential decay law at the end of each iteration. Initially, the social network is a random graph of degree K , and each social connection has the same decay time. When two agents perform activities at the same location, two situations can occur: (a) if there is a social connection between them, the strength of the connection is reinforced and its decay time increases and (b) if there is no social connection, a new social connection is created. This task requires $O(AK)$ operations for the enumeration of the social links and $O(C)$ operations to browse the cells for detecting interactions between agents. However, the detection of the existence of a social link between two given agents would require $O(AK)$ operations for each cell to browse the connections. An alternative solution would be to store the connections of a given agent in a hash table. Both solutions are costly in memory or computation time. For those reasons, a slightly different implementation, which is equivalent statistically, is

adopted. Every connection has the same decay time, but multiple redundant connections can exist between two given agents. When an agent visits a particular cell from his or her own elite buffer, the cell-agent pair is stored until a second agent visits the cell. A new social connection is then created between the two visitors. After one iteration, the result is a larger number of connections than at the beginning. They are of equal strength but with potential redundancy (a given pair of agents can appear several times). Identifying the redundancy would be costly. Instead, the total connection strength is kept constant in the system by deleting connections randomly so that the total number of connections is constant.

IMPLEMENTATION ISSUES

The code of the simulation has been written in Java. Input and output files use the XML file format, which is well suited to variable length content such as the description of the individual plans with multiple stops. The goal being to simulate 10^6 agents on a single CPU, some performance concerns have to be taken into account.

Cells

Each grid cell is stored as an individual object. With 10^5 cells, it is not crucial to store the cell attributes as plain arrays. This allows keeping cell characteristics private and having cell references. Because cells have to be compared often during the learning process to determine whether a new cell is already known to an agent, it is far more efficient to use the identity operator \equiv than the default Java `equals()` method, which is the equivalence operator. That is valid as long as the cells are not dynamically allocated or cloned once the simulation starts. Cells have to keep references to agents that visit them (see the socialization stage). The average number of visits per cell is small, $O(SA/C)$, so that the overhead of a dynamic container (e.g., vector) could be allowed for. However, in the simple pairwise interaction described above, only the information about the last visitor is needed. Therefore, it is sufficient to maintain a hash table of visited cells and visitors.

Agents

It is tempting to have a dynamic container steadily increasing in size for the agent memory. However, that would completely ruin the performances and is not compatible with the assumption that only a limited number of simultaneous options can be memorized. Each cell is referenced on average by $O(AN/C)$ agents so that the initial coverage is sufficient to ensure that there is not any information missing about the environment.

Random Numbers

A typical bottleneck of this kind of simulation is the computation of random numbers. A priori, $O(AK)$ random numbers have to be computed for each single iteration of the information exchange stage. That quickly becomes prohibitive and can be avoided by using two integer random seeds at the beginning of the exchange stage. One is used to pick a cell from the informer agent, the other to replace a cell in the memory of the informed agent. These two pointers can be simply incremented from one social link to the other because no correlation exists between social links and they are accessed in an a priori random order.

Update Sequence

The original ordering of the agents is that of the XML source and could be biased. Agents are processed sequentially during the evaluation phase: each agent uses one of his or her elite plans and marks the visited cells. Therefore, the probability of creating a social connection between agents decreases with their distance in the list. To avoid that potential bias, the list of agents is randomized at the begin-

ning. Social connections are also stored initially in a random list because they are processed sequentially during the learning phase. New social connections are created at the end of the list. The decay process rearranges them randomly.

RESULTS

The simulation is tested on a real-world example for the Zurich region for which a high-resolution transportation network and a land-use raster are available (Figure 2). The area covers an approximately 50×50 -km² area, in which about one million inhabitants are living. The land-use utility values R_i are generated on the basis of census data. Random plans with one or two intermediary stops are generated for 10^6 agents who are distributed in the area according to job and housing densities. The home-to-work pairs are computed using an external model written by one of the authors [see Marchal for computation of rent values presented in Figure 2 (9)]. The initial social network that connects them is a random graph. Obviously, this is not realistic, but the intention is to evaluate only the computational feasibility in this preliminary work.

Figure 3 presents the evolution of the sum of the scores of all agents during the iterative process. It can be seen that the process converges in a few dozen iterations but that the choices are not optimal because the utility does not reach the maximum value obtained with a full enumeration of the alternatives. That is because some information is lost in the process. It can be recovered slowly through the exploration. However, some information can still be lost forever because the exploration mechanism is local: the agents cannot jump to explore a totally new area. This limitation could be removed easily. Still, the value of the plateau is high enough (more than 90% of the maximum utility) to ensure that plausible strategies have been selected. It remains to be studied how much this is compatible with

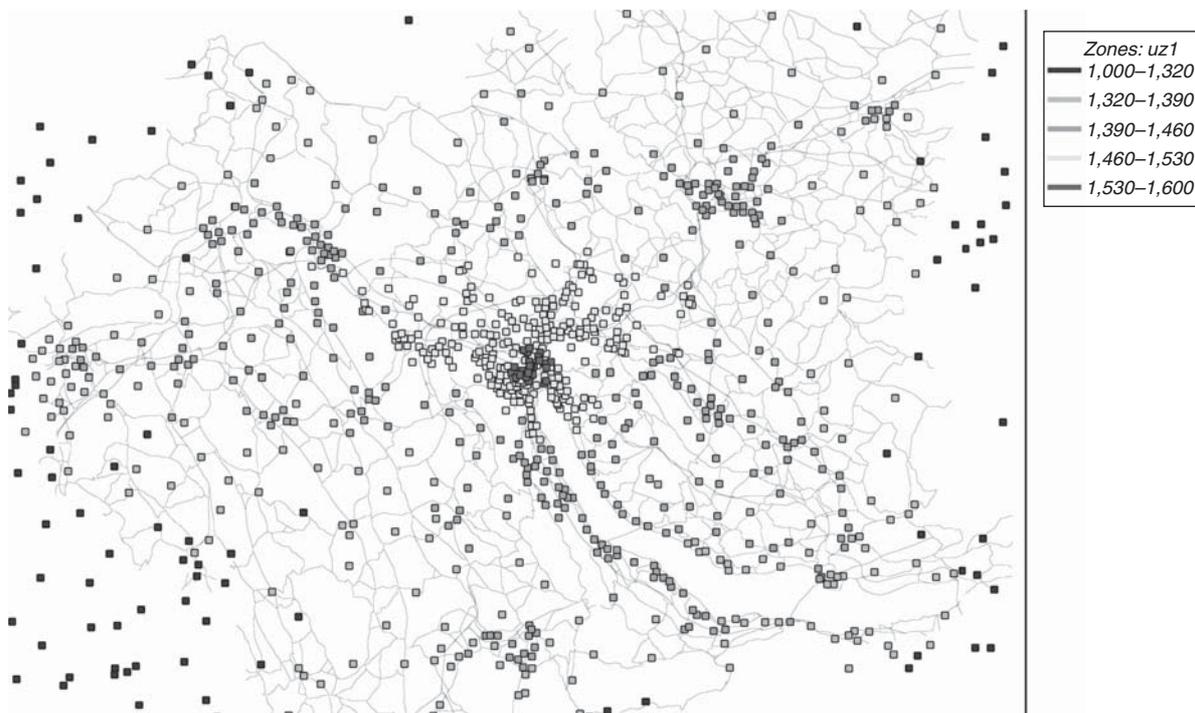


FIGURE 2 Zürich area: transportation network and rent values. (User: Mon., July 3, 2003; database: ZH_LU; network-Network 5. METROPOLIS 1.4-.) (Copyright © André de Palma 2001–2003. All rights reserved.)

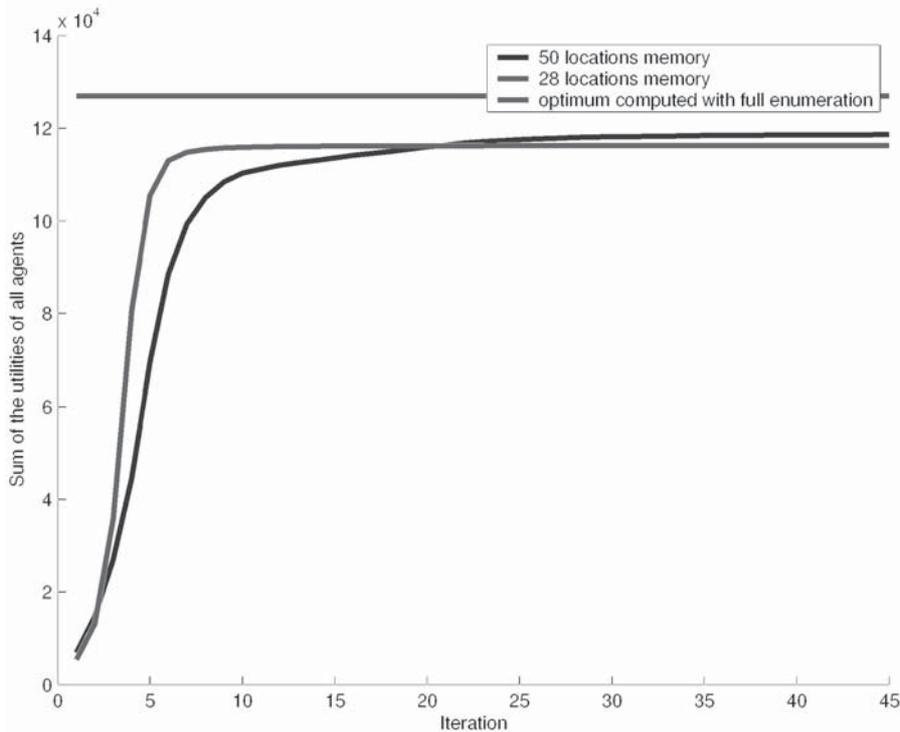


FIGURE 3 Performance of multiagent simulation.

empirical evidence. Note that the size of the memory of the agent affects the convergence properties only slightly.

Figure 4 illustrates a typical spatial adaptation process for a single agent with a two-stop plan: home-work-leisure-shopping-home. On the first iteration, leisure ($L1$) and shopping ($S1$) are performed at the home location because the agent ignores good locations to perform these activities. On the second iteration, the agent learns that location $L2-S2$ is good for one of the two activities, hence making the extra trip distance worth it. On the third iteration, the agent discovers that the area around $S3$ has a high utility for shopping. During the fourth and fifth iterations, the agent keeps shopping close to that area and optimizes only the leisure location ($L3 \rightarrow L4 \rightarrow L5$).

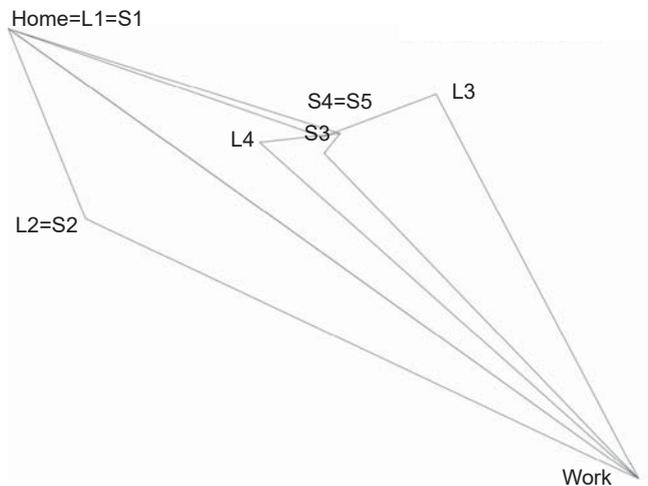


FIGURE 4 Adaptation of location for two-stop plan.

Figure 5 shows the evolution of the distribution of social connections in the system. The initial condition is a random graph of degree $K = 20$. On the first iteration, the distribution is roughly an $N(20,1)$ distribution. Progressively, the distributions shift to the left. Ultimately, there is a self-sustained distribution that can be approximated by an $N(13,4)$ distribution. Therefore, most agents maintain between nine and 17 connections. Note that redundant connections between a given pair of agents are not identified and are therefore counted multiple times. The decrease (from 20 to 13) in the average number of social connections indicates that the spatial interaction is not sufficient, in this case, to sustain 20 connections on average. That is dependent on the properties of the land-use data such as the concentration of high-utility areas. It also depends on the total number of agents in the system. (The area under the curve is equal to the total number of connections and thus is constant.)

All the experiments were done on a computer equipped with an Intel Pentium 4 clocked at 2.5 GHz. The typical simulation performance for 100 iterations of a system with 10^6 agents is below 1 h of central processing unit time. That is for plans that have only one or two intermediary stops. In regard to memory requirement, about 400 MB of RAM are needed. Obviously, the simulation of larger systems and more sophisticated plans with more than two stops will require distributing the workload over several computers. In particular, extending the framework to integrate other travel decisions, such as the timing of the different trips, is desired.

CONCLUSIONS AND PERSPECTIVES

Many aspects still need to be addressed to improve the realism of the model. In particular, validating the model using results of a recent survey conducted in Germany and Switzerland is planned (10).

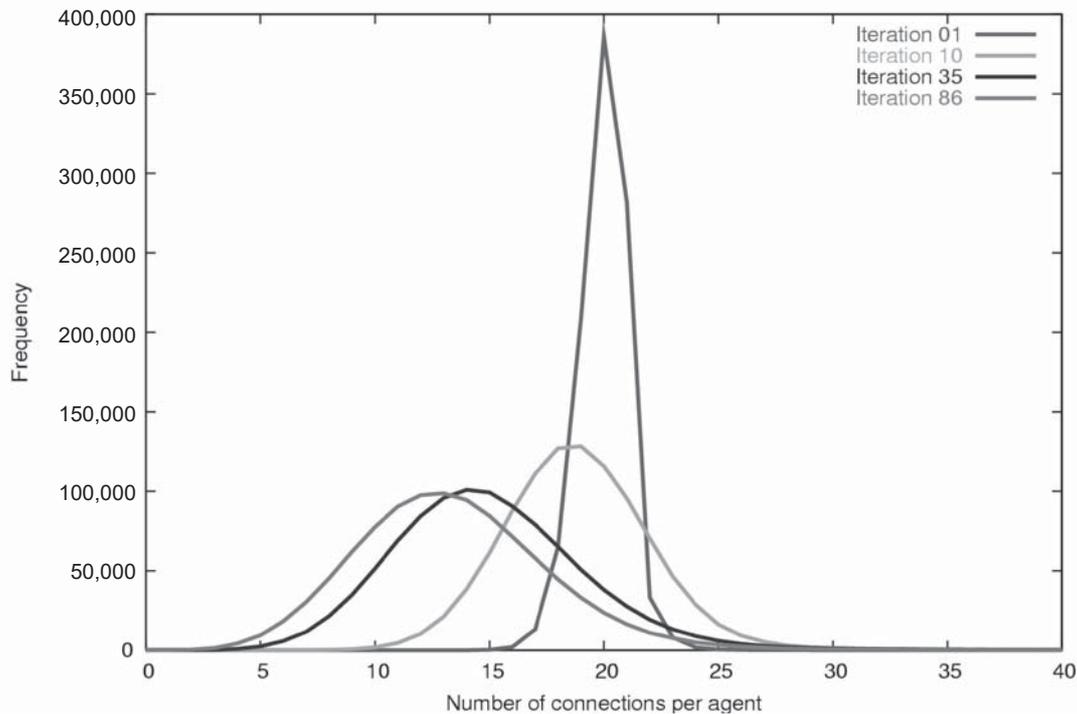


FIGURE 5 Distribution of social connections.

The goal of this paper was to study the factors that influence leisure trips. Respondents were asked to give the location and the frequency of visits to their five closest friends or relatives. Empirical results suggest that the purpose of an important share of leisure trips is to visit social connections and that the number of known locations is small. These data should, in principle, allow validation of the model and a comparison of the spatial distribution of social connections with that of this model. This preliminary work has shown that a multiagent-based approach to the location of secondary activities is technically feasible and behaviorally plausible for high-resolution data sets. The fact that agents cooperate in even some simplistic way yields an important gain in regard to computation workload. That has to be compared with the standard practice in transportation science in which, typically, it is assumed that users are in the situation of a noncooperative Nash equilibrium.

ACKNOWLEDGMENTS

The authors thank K. Axhausen and PTV AG for providing them with the data for the Zurich area. The authors acknowledge the extremely helpful comments of the anonymous referees on the earlier version of this paper. Credits go also to A. Altenhoff, who wrote a first implementation of the simulation. Computer resources for the simulations were provided by the Computational Laboratory (CoLab) at ETHZ.

REFERENCES

1. P. Waddell, A. Borning, M. Noth, N. Freier, M. Becke, and G. Ulfarsson. Microsimulation of Urban Development and Location Choices: Design

- and Implementation of UrbanSim. *Networks and Spatial Economics*, Vol. 3, No. 1, 2003, pp. 43–67.
2. Axhausen, K. W., A. Zimmermann, S. Schönfelder, G. Rindsfuser, and T. Haupt. Observing the Rhythms of Daily Life: A Six-Week Travel Diary. *Transportation*, Vol. 29, No. 2, 2002, pp. 95–124.
3. Raney, B., N. Cetin, A. Völlmy, M. Vrtic, K. Axhausen, and K. Nagel. An Agent-Based Microsimulation Model of Swiss Travel: First Results. *Networks and Spatial Economics*, Vol. 3, No. 1, 2003, pp. 23–41.
4. De Palma, A., and F. Marchal. Real Case Applications of the Fully Dynamic METROPOLIS Tool-Box: An Advocacy for Large-Scale Mesoscopic Transportation Systems. *Networks and Spatial Economics*, Vol. 2, No. 4, 2002, pp. 347–369.
5. Ben-Akiva, M., and S. R. Lerman. *Discrete Choice Analysis*. MIT Press, Cambridge, Mass., 1985.
6. Domencich, T. A., and D. McFadden. Urban Travel Demand. In *Urban Travel Demand*, No. 93 in *Contributions to Economic Analysis* (D. W. Jorgenson and J. Waelbroeck, eds.), North-Holland Publishing, Amsterdam, Netherlands, and American Elsevier, 1975.
7. Arentze, T., and H. J. P. Timmermans. Representing Mental Maps and Cognitive Learning in Micro-Simulation Models of Activity-Travel Choice Dynamics. In *Proc., Meeting of the International Association for Travel Behavior Research*, Lucerne, Switzerland, 2003. www.ivt.baum.ethz.ch/allgemein/iatbr2003.html. Accessed June 2004.
8. Bonabeau, E., M. Dorigo, and G. Theraulaz. *Swarm Intelligence: From Natural to Artificial Systems*. Santa Fe Institute Studies on the Sciences of Complexity. Oxford University Press, 1999.
9. Marchal, F. A Trip Generation Method for Time-Dependent Large-Scale Simulations of Transport and Land-Use. *Networks and Spatial Economics*, forthcoming, 2005.
10. Schlich, R., A. Simma, and K. W. Axhausen. Determinanten des Freizeitverkehrs—Modellierung und empirische Befunde. *Arbeitsberichte Verkehrs—und Raumplanung 190*, Institut für Verkehrsplanung und Transportsysteme (IVT). ETH Zürich, Switzerland, 2003. www.ivt.baug.ethz.ch. Accessed June 2004.

The Artificial Intelligence and Advanced Computing Applications Committee sponsored publication of this paper.